AI Magazine Volume 6 Number 2 (1985) (© AAAI)

TECHNOLOGY TRANSFER

Developing A Knowledge Engineering Capability in the TRW Defense Systems Group

Edward C. Taylor

TRW Defense Systems Group, Redondo Beach, CA 90278

The TRW Defense Systems Group develops large man-machine networks that solve problems for government agencies. Until a few years ago these networks were either tightly-coupled humans loosely supported by machines—like our ballistic missile system engineering organization, which provides technical advice to the Air Force, or tightly-coupled machines loosely controlled by humans—like the ground station for the NASA Tracking and Data Relay Satellite System. Because we have been producing first-of-a kind systems like these since the early 1950s, we consider ourselves leaders in the social art of assembling effective teams of diverse experts, and in the engineering art of conceiving and developing networks of interacting machines But in the mid-1970s we began building systems in which humans and machines must be tightly coupled to each other—systems like the Sensor Data Fusion Center (Figure 1). Then we found that our well-worked system development techniques did not completely apply, and that our system engineering handbook needed a new chapter on communication between people and machines We're still writing that chapter, and it won't be finished until we can add some not-yet understood cognitive psychology and some not-yet fully developed artificial intelligence techniques Nevertheless, we have learned some lessons worth passing along.

Sensor Exploitation Problem

In the Sensor Data Fusion Center concept (see Figure 1), diverse sensors on various kinds of platforms monitor enemy operations beyond the horizon, and the sensor reports travel over a communication system to a correlation processor. The processor automatically merges reports and stores descriptions of each detected item Operators query the files constructed by the correlation processor, and the displays respond in formats selected by the operators The operators use these displays to deduce the nature and location of significant enemy deployments By monitoring changes in the enemy deployments they try to anticipate enemy initiatives so they can help their commander counter them.

The fusion system did what it was designed to do: It demonstrated the correlation, query, and display capabilities in a test bed in 1981, and the European Command, U.S. Forces has deployed the system in a limited-operational capability mode. What has yet to be demonstrated is that these processes enable analysts to help a commander outwit an enemy. The crucial question is no longer whether sensor reports can be rapidly correlated, but rather how well humans can sort through large amounts of correlated sensor data to assess situations rapidly and accurately. Because this question pertains to many of the systems built by the TRW Defense Systems Group, we started a separate program of experiments in military situation assessment in 1980

Some Man-Machine Experiments

In one set of experiments, conceived of and managed by Daniel Newell of TRW Defense Systems Group, a simulator produced the sensor outputs that might accompany a Warsaw Pact attack on a NATO member. We fused the outputs and made them available on displays controlled by experienced tactical intelligence analysts. We briefed the analysts on the known situation before the attack, and told them to watch for and assess any important developments. The analysts could monitor current sensor reports, call up any past sensor reports, and task any sensor The sim-

This story would not be worth telling if it were not for the government people who took the crucial step of testing our systems with experienced military operators It is a pleasure to acknowledge the overall guidance and support for development of the demonstration work station by Daniel Wiener of the Joint Tactical Fusion Program Office Daniel Ventimiglia of the Rome Air Development Center provided guidance and support for the work station demonstration, and Captain Richard Radcliffe of the Joint Tactical Fusion Program Office provided the development and testing of the correlation tuner

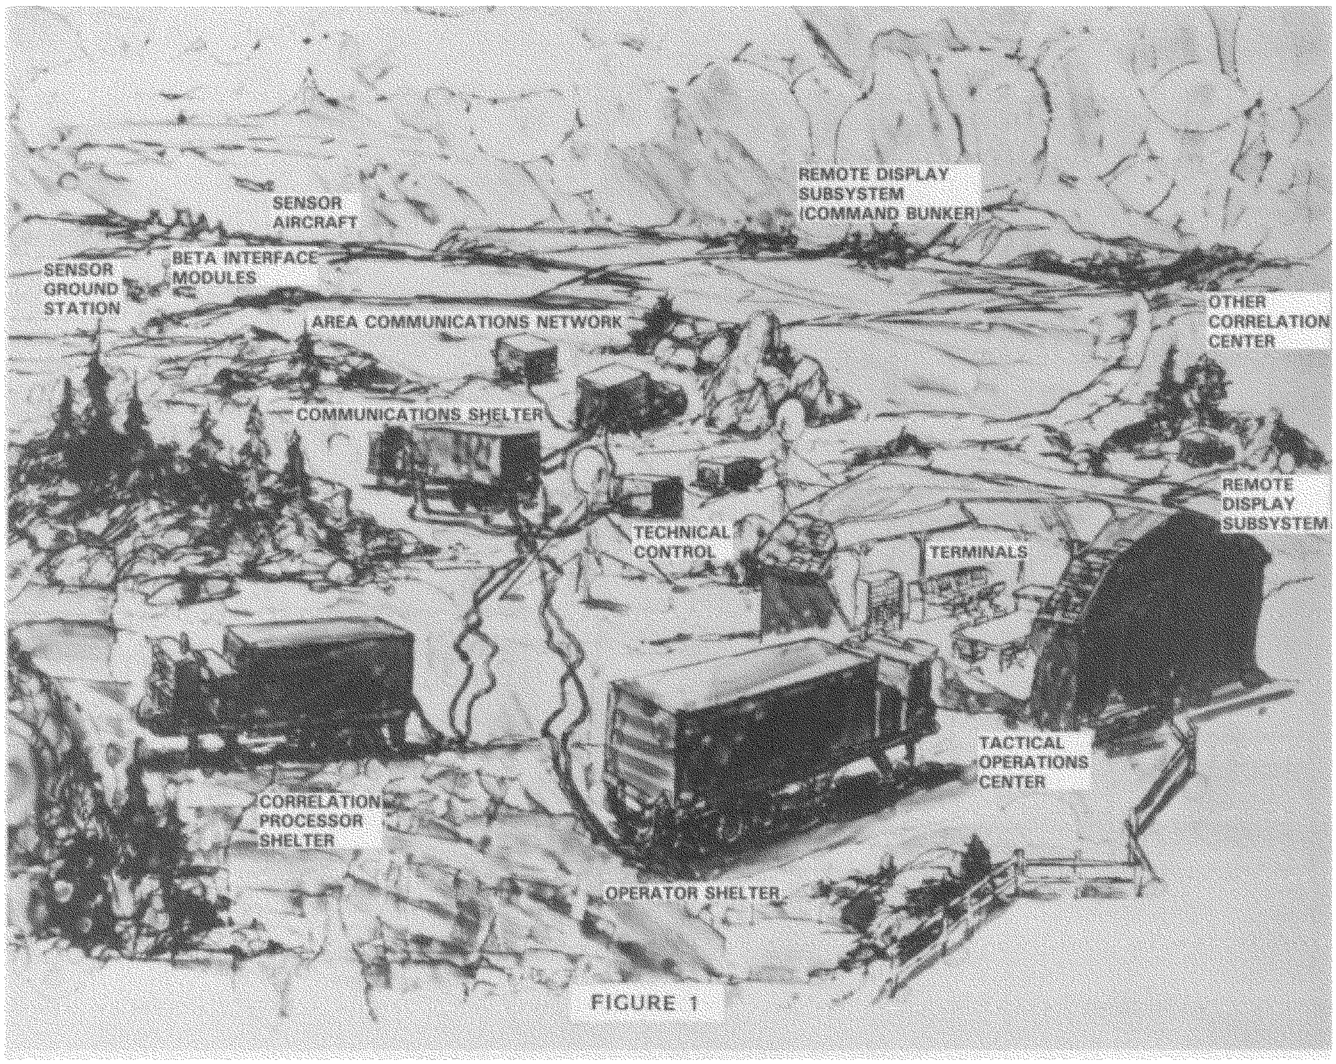

Figure 1.

ulated attack developed, and after an hour we asked the analysts for a full situation assessment. We made several observations:

- It was difficult for most analysts to abstract good situation assessments from such large data streams.
- The approach of each analyst to situation assessment was different.
- The top analysts were good situation assessors but not fast enough.

These observations surprised us, because we initially believed that the crucial tactical intelligence problem was to supply enough sensor information to analysts. The observations suggest that the problem is not limited information, but limited attention, and that the natural cognitive limits of humans are the crucial constraint in military situation assessment.

Cognitive limits affect each person differently, and our experiments suggest that these differences cause wide vari-

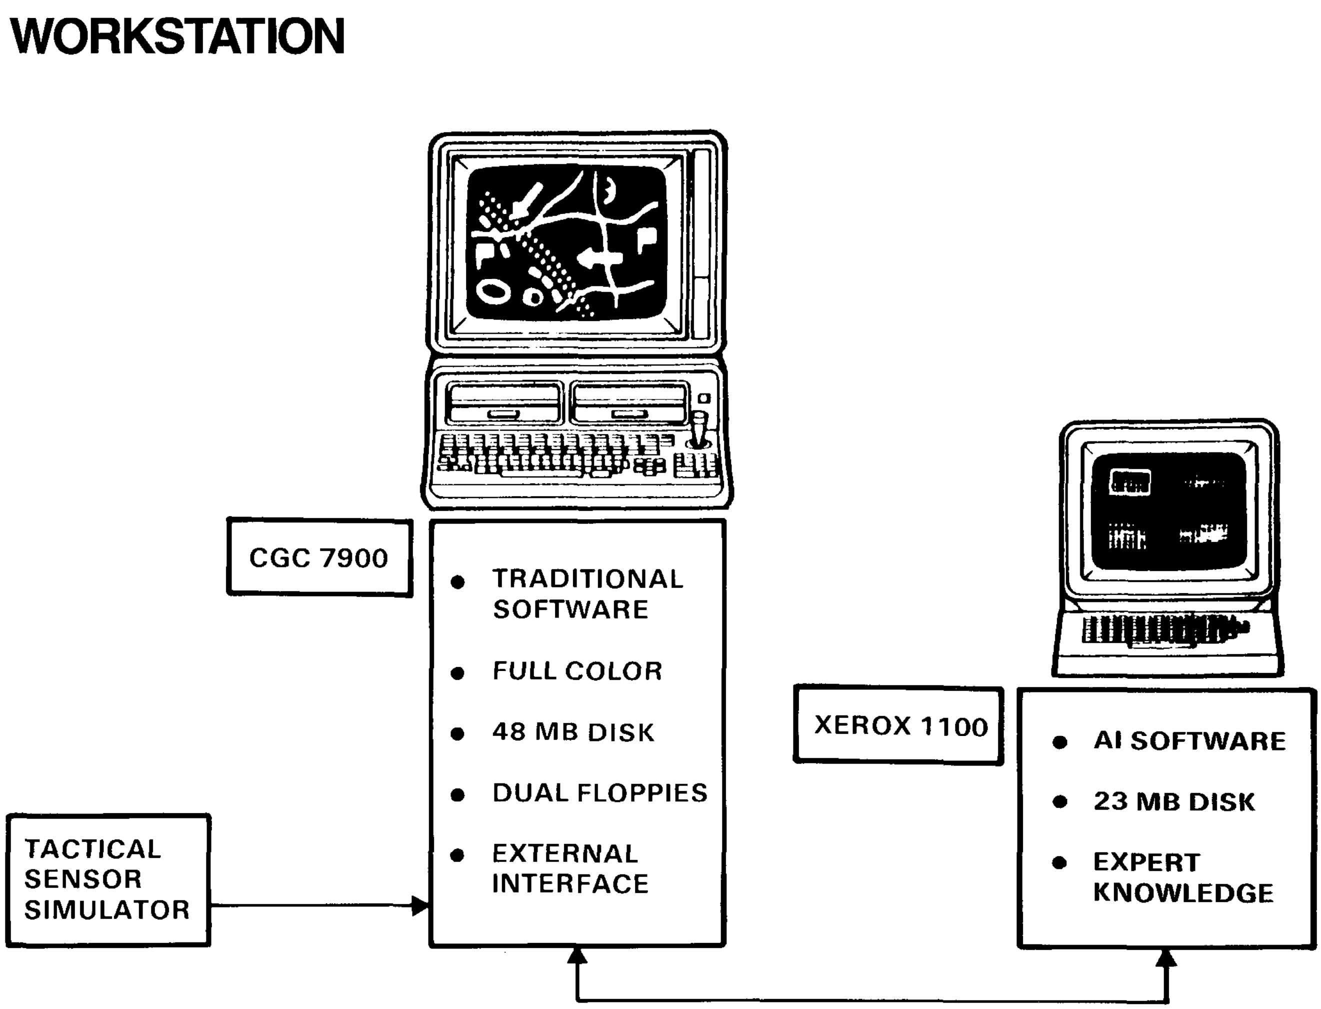

Tactical Intelligence Workstation Experiment.

Figure 2.

ations in the quality of the performance of real time intelligence systems. The same is presumably true of all other command and control system elements that depend on human skill. This problem will need attention if we are to continue exploiting the U.S. technological advantage in military systems.

The promising observation was that there are good analysts who, if they have enough time, do well at situation assessment because they seem to know what to look for. As in past experiments with chess players, computer programmers, and physicists, the good problem solvers and bad problem solvers think at about the same speeds, and examine about the same number of items. But the good problem solvers tend to get better ideas to think about because they have better and more organized knowledge We realized that if we could transfer knowledge from the good analysts to the sensor-data displays, and design the displays to prompt analysts at the appropriate times, then we could improve the ability of analysts to assess situations rapidly. Our long-term objective became to develop and demonstrate a generic knowledge-based system that represents, organizes, and monitors reports about changing situations; generates and evaluates possible situation interpretations; and interacts with operators in helpful ways.

Testing A Conceptual Solution

To develop a prototype situation assessment system, we obtained ROSIE (Rule-Oriented System for Incorporating Expertise), a LISP-based language developed for DARPA by the Rand Corporation, and adapted it for incorporation into a system that now operates on a Xerox 1100 LISP Machine. ROSIE is an English-like language for building accessible, intelligible, and malleable knowledge bases. Our first system, completed in 1982, was a situation assessment aid for space defense. It used knowledge about satellite and booster phenomena and access to our library of trajectory calculation programs to develop and check hypotheses about newly discovered objects in space.

The performance of our first situation assessment system was promising, and the Air Force agreed to try it out with practicing tactical intelligence analysts. We expanded the system, and it became a demonstration model for a tactical intelligence work station. The demonstration work station was an integrated set of appropriately modified off-the-shelf hardware and software packages (see Figure 2). The hardware includes a Xerox 1100 DOLPHIN, which executes the ROSIE-based situation assessment system developed for space defense, but with the space defense knowledge replaced by a tactical intelligence analyst's knowledge. The hardware also includes a Chromatics CGC 7900 graphics display system. Displays are developed and executed by TRW's Operator-Machine Interface Module, which provides facilities for manipulation of maps, the military, symbology, annotation, and distance computations. The CGC 7900 also performs the special processing associated with the reception and fusion of sensor reports.

The major problem in developing the demonstration work station was the selection and implementation of an appropriate subset of the knowledge required to understand sensor reports in an air-land battle. The knowledge subset had to be authentic enough to satisfy practicing military intelligence analysts that they were not dealing with a toy; it had to help interpret a long, connected series of significant events to support a demonstration lasting several hours; and it had to fit on the DOLPHIN despite the many supporting definitions required for even the simplest piece of combat knowledge. The designer, Dan Snell, a former intelligence analyst (now a TRW system engineer), correctly guessed that the knowledge required to locate second-echelon divisions halting in assembly areas satisfied these conditions.

The demonstration work station receives simulated reports from a variety of sensor syste The CGC 7900 performs data reduction functions to produce a display of combat nodes on the battlefield. The DOLPHIN accesses this information, along with reports that describe events. The DOLPHIN contains rules about how and why enemy forces operate, and it helps the analyst understand what is happening.

A typical interaction might go as follows (paraphrased for clarity):

Work station	I think the second echelon force is in a long halt, and that it will be stationary for four hours.
Operator	How did you figure that out?
Work station:	I have the enemy doctrine you taught me, and I have relevant sensor reports It looks like the second echelon has gone off-road, and at this stage they would only go off-road to rest and reassemble As you taught me, when they go off the road, they rest for four hours
Operator	Yes, but from what I've been seeing over the past twenty-four hours, I've changed my mind—I think they'll only rest for two hours
Work station	Then your commander has one hour to go if he's going to attack them, because they've been resting for one hour already

While the analyst might or might not reach a similar conclusion without aid, the work station helps by sorting through a morass of data and making tactically significant associations in a stream of reports that would be very difficult to keep up with in the heat of battle.

We tested the demonstration work station with eleven analysts brought in from the field by the Air Force. Before the test, they were briefed on the system concept and the nature of the demonstration, and given a day of hands-on training. The briefings noted that the system is a prototype, that the knowledge in the system is limited to certain aspects of second-echelon behavior; and that the system is slow. After operating the work station for several hours of simulated warfare, the analysts reported their opinions. They were able to see through the slowness of the system and the limits of its knowledge. They recognized that the hypothesizing, arguing, and explaining provided by the work station is just what they need to help them deal quickly and effectively with large amounts of data. (One analyst said, "When it does something, it does what I would do if I had the time to do it.") They left us with a list of recommendations for turning the prototype into an operational system Chief among these suggestions were more capacity for the knowledge they would want to put in themselves, and more speed. But there was no question that the concept of a knowledge-based work station is a significant step toward improving the ability of humans to handle large amounts of sensor data

The situation assessment workstation activities have now divided into two parts: field tests in Europe, and laboratory research at Rome Air Development Center. Daniel Weiner of the Joint Tactical Fusion Program Office directs the field test activity, and he is directing the development of an advanced version of the workstation for experimental use in military exercises. Daniel Ventimiglia at the Rome Air Development Center directs the laboratory research, and he is directing the development of improvements recommended by the analysts in the prototype tests.

Other Knowledge Engineering Developments

While we were experimenting with the incorporation of expert knowledge into machines to aid intelligence analysts, we realized that other kinds of knowledge based-systems could be immediately useful, and would be easier to build. Supporting fusion system operations is a good example. The quality of the performance of the fusion system depends on the quality of the settings of several thousand correlation parameters in the algorithms that fuse sensor reports. Proper setting of these parameters requires knowledge of the observables, knowledge of the sensors, and knowledge of the system design. One of the people who has the required knowledge—Charles Zumba (now at SAI, Inc.)—was working in the TRW Defense Systems Group on the design of the fusion system at the time we were learning about knowledge-based syste He pointed out that a knowledge-based tuner would improve performance, and he agreed to participate in the design. We built a demonstrator with our own funds, and on the basis of the demonstration the Joint Tactical Fusion Program Office concluded that a scaled-up version of the knowledge in the tuner would enable it to improve the operation of sensor-data fusion systems in the field Daniel Wiener directed the development of the larger knowledge base and Captain Richard Radcliffe conducted the field demonstration with it. The demonstration showed that the tuner improves operator understanding of the fusion algorithms, and that the use of the tuner improves the performance of the system. As a result, we are now on contract to deliver a knowledge-based tuner for permanent use at various sites in the United States and in Europe.

We have identified a number of other problems that might be solved by artificial intelligence techniques. Our development approach is first to build a demonstrator with company funds, then to build a prototype on government funds if the demonstration succeeds, and finally to use the results of prototype testing to specify a system for field use We believe that the development of the systems should be led by the people close to the application, and not by a central organization of artificial intelligence specialists. We have at least ten organizations that have produced demonstration systems, and most of them have obtained funding from customers to produce prototypes. (The artificial tuner described above is the first system to get government funding to be put into field use)

Current Status

We started our knowledge engineering program with the situation assessment experiments in 1980, and began building knowledge-based systems in 1981. The idea of using knowledge-based systems to support the operation and maintenance of our large systems caught on and spread. Now we have about 40 people distributed across ten knowledge engineering centers at six geographical locations. At first the efforts were entirely supported with special long-term research funds, but today they are all supported by customers and program managers with discretionary funds. So far, we have delivered each product (three deliveries, so far, several more to come) on time, within budget, and with promised performance. We think we have made an effective transfer of knowledge-engineering technology from the academic world to the aerospace industry.

Lessons Learned

In the process of building our knowledge engineering capabilities, we have learned lessons not found in the expert system literature The most important lesson is that expert systems are best conceived of by the experts themselves. The idea of a team of knowledge engineers who identify human tasks that can be profitably automated did not work here (We think mastery of the domain knowledge is required to make good decisions about building expert systems, and the mastering of significant domain knowledge takes years.) What did work was company-wide briefings that spelled out what might be done with this new technology, and that challenged experts in the company to propose and participate in the development of useful expert systems. Our successful systems were conceived of by people who had never written a line of LISP.

We also found that the people who are good at building knowledge-based systems (once the systems have been defined by domain experts) have two things in common They are curious—they read omnivorously and voraciously—and they are clever programmers. They typically have unconventional educational backgrounds. One has a Ph.D. in philosophy; another a Ph.D. in linguistics; two have M.A 's in political science; and two have no degrees at all. Finding such people and exposing them to a good consultant (Our good consultant was Frederick Hayes-Roth, now of Teknowledge) gets a competent knowledge engineering center going. We are no longer concerned about the shortage of people highly trained in artificial intelligence, but we are still concerned about the abiding shortage of imaginative people.

Finally, our attempt to transfer tasks from humans to machines has given us a new respect for human capabilities Figure 3 shows our view of the command and control problem. Sensor systems produce data in alphanumeric form, but only human staffs can convert these data into a situation assessment. A commander thinks about this situation assessment until he gets an idea about what he

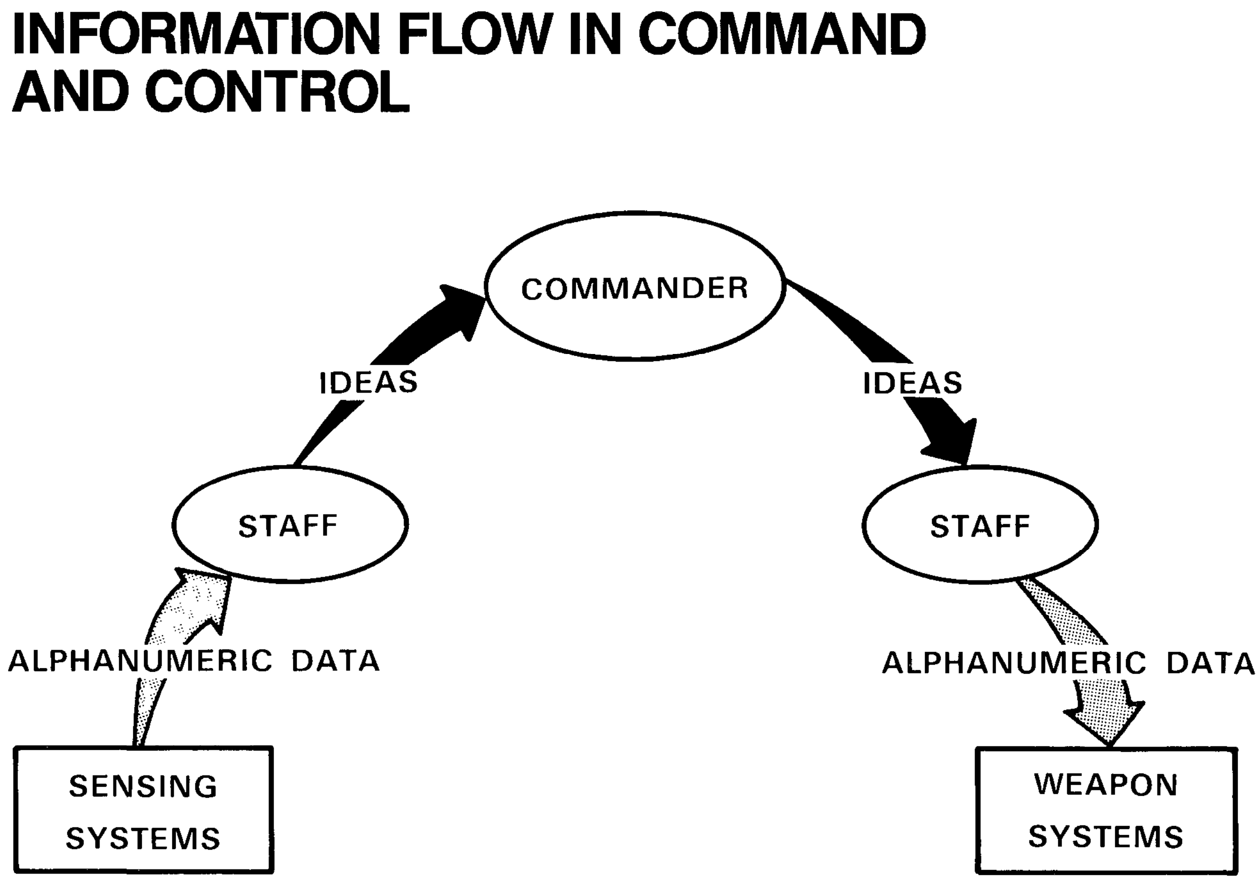

Figure 3.

wants to do. Then he presents his ideas to his staff, and the staff converts these ideas into an alphanumeric battle plan that tells the weapon systems what to do. We are good at making systems that manipulate alphanumeric data, and humans are good at manipulating ideas.

The problem is the rapid conversion of alphanumeric data to ideas, and ideas to alphanumeric data. Only humans can make these conversions, but our experiments (and many other experiments) show that humans can't convert fast enough to help commanders exploit our sophisticated sensors and weapons in the rapid-paced wars we are preparing to fight. We have to invent knowledge-based machines that enable staffs to do this otherwise undoable job. We think our work station demonstration shows such a solution exists.

SPECIAL ISSUE ON AI IN MANUFACTURING

The Winter issue (Volume 6, Number 4) of the the *AI Magazine* will focus on a single topic: AI in Manufacturing. Mark Fox (Carnegie-Mellon University) will serve as Guest Editor of that issue. One of the features will be a snapshot of the people, organizations and activities at the intersection of AI and Manufacturing, circa 1985 A primary source of data for that picture is our own Association. In the previous issue of your Magazine you will find a tear-out survey form. This survey is sponsored by the AAAI. The Association will have custody of the survey forms, and the data will be available to its members. After the survey is completed, you can request copies from the AAAI office, 445 Burgess Drive, Menlo Park, CA 94025.

Most of us hate filling out forms, so we've tried to make it as quick and painless as possible to extract the information needed. The results of the survey will be published in the special issue.

Thank you for your cooperation.